BMDO

# REPORT DOCUMENTATION PAGE

AFRL-SR-BL-TR-00-

0047



| 1. REPORT DATE (DD-MM-YYYY) | 2. REPORT TYPE | 3. DATES COVERED (From - To) |
|---|---|---|
| 03-01-2000 | Final Technical | August 1998 to September 1999 |

**4. TITLE AND SUBTITLE**
Optical Dynamic RAM

**5a. CONTRACT NUMBER**
F49620-98-C-0052

**5b. GRANT NUMBER**

**5c. PROGRAM ELEMENT NUMBER**

**5d. PROJECT NUMBER**

**5e. TASK NUMBER**

**5f. WORK UNIT NUMBER**

**6. AUTHOR(S)**
Johnson, Alan E.
Maniloff, Eric S.
Mossberg, Thomas W.

**7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)**
Templex Technology
1850 Millrace Drive, #1
Eugene, OR 97403

(541) 683-7474

Cage Code: 06RG3

**8. PERFORMING ORGANIZATION REPORT NUMBER**

**9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)**
USAF, AFMC
Air Force Office of Scientific Research
801 N. Randolph Street, Room 732
Arlington, VA 22203-1977 (703) 696-8573
(BMDO 97-011)

Cage Code: FA9550

**10. SPONSOR/MONITOR'S ACRONYM(S)**
BMDO/AFOSR

**11. SPONSOR/MONITOR'S REPORT NUMBER(S)**

**12. DISTRIBUTION / AVAILABILITY STATEMENT**
Distribution Statement A: Approved for public release. Distribution is unlimited.

**13. SUPPLEMENTARY NOTES**

**14. ABSTRACT** Optical Dynamic RAM (ODRAM) is a memory architecture designed for fast access to large stores of data. In particular, ODRAM would enable microsecond random access to data stores of up to fifty gigabytes in a single unit. This is to be compared to the millisecond random access time achievable with existing disk-based technology. Relative to ODRAM, the only common technology that is faster is silicon RAM. During Phase II, we have designed, characterized and demonstrated several of the subsystems and data storage techniques required for the construction of ODRAM. Concurrently with our development efforts, the price of the competition, silicon RAM, has fallen dramatically. In fact, it is now possible to achieve nanosecond random access times to comparably large data stores using silicon DRAM at a price equal to our projected ODRAM price. Given the advantage in speed and the current commercial popularity of silicon DRAM, the chances for commercial success with ODRAM are vanishingly small.

**15. SUBJECT TERMS**
Optical Memory; ODRAM

| 16. SECURITY CLASSIFICATION OF: | | | 17. LIMITATION OF ABSTRACT | 18. NUMBER OF PAGES | 19a. NAME OF RESPONSIBLE PERSON |
|---|---|---|---|---|---|
| a. REPORT | b. ABSTRACT | c. THIS PAGE | | | Alan E. Johnson, PI |
| Unclassified | Unclassified | Unclassified | Unclassified | 17 | 19b. TELEPHONE NUMBER (include area code) (541) 683-7474 |

## Project Overview:

The primary objective for this project was the development of an optical memory technology to be brought to market in the near term following completion of the Phase II contract. This memory technology would enable very rapid random access (10 microsecond latency) to large (tens of gigabytes) quantities of data. With a substantial speed advantage over magnetic hard drives and other mechanically accessed storage media, our major competition in the marketplace would be semiconductor DRAM.

During Phase I, we defined memory architecture and performance targets. Our initial (year 1) Phase II objectives included high multiplicity single spot data storage, multispot addressing and construction of a first generation prototype, ODRAM-1. Upon trying to achieve high multiplicity single spot storage, we found that the maximal single spot capacity was several orders of magnitude below that expected from initial estimates based on the homogeneous linewidth measured in the low excitation limit. Rather than accept this limitation and construct ODRAM-1 with a significantly reduced level of performance, we modified our first year objectives to emphasize development of new data storage techniques that would enable fuller utilization of the theoretical material storage capacity. To that end, we demonstrated a new technique, Spatially Distributed Spectral Memory, which enabled us to regain a substantial fraction of the lost capacity.

Concurrently to our development efforts, the price of semiconductor DRAM continued to drop at a historically unprecedented rate. In the past four years, the price of semiconductor DRAM has fallen over two orders of magnitude, plummeting to a low value of $500/Gigabyte, which is close to our initial price target for ODRAM ($25,000 for several 10's of gigabytes of capacity). Given the likelihood of continuing price reduction in a competitive technology with access times substantially lower than ODRAM, the remaining technical hurdles in ODRAM development, and the difficulties associated with introducing a new technology when there is an entrenched alternative, the probability of developing ODRAM as a viable commercial product is extremely low.

## Introduction to ODRAM

### 1) The promise of spectral memories

Persistent spectral hole burning in frequency selective materials has long held promise for implementation of high capacity optical memories by enabling the storage of multiple bits in a single storage location. In spectral memories the theoretical storage capacity (number of bits) of a single storage location is equal to the ratio of the inhomogeneous linewidth to the homogeneous linewidth ($\Gamma_i/\Gamma_h$) and can be as high as ten million in some materials.[1,2] In the earliest proposed implementations, each and every bit would be assigned a specific color channel as well as a specific location on a storage material. In other words, the data would be stored in three dimensions, one spectral and two spatial. (Note that generalizations to four dimensions, three spatial and one spectral, are also possible). Beam steering techniques and/or motion of the storage material would be used

to access the desired spatial address and the frequency of a tunable laser would be used to access the desired spectral address.

## 2) Time-domain approaches to spectral data storage

A potential difficulty in the frequency domain implementation of spectral memories as described above arises upon trying to access spectral channels with very narrow homogeneous linewidths. In particular, the Fourier relationship between time and frequency requires that the minimum interrogation time required for access to a particular channel is the reciprocal of the channel width. For storage materials with sub-MHz lines, this limit requires maximal single channel bit rates substantially less than 1 Mbit/s. Time-domain and swept-carrier approaches to spectral storage using frequency selective materials have been demonstrated to overcome the time-bandwidth limitations (input/output speed) encountered when directly accessing spectral channels.[2-7] In the time-domain approach to spectral storage, the storage material is first illuminated with a reference pulse and is subsequently illuminated with a data sequence. The spectral interference pattern between the Fourier spectrum of the reference pulse and the Fourier spectrum of the temporal data sequence is stored in the frequency selective material. Subsequent illumination of the storage material with a read pulse that is substantially similar to the reference pulse causes the generation of an outgoing replica of the stored data sequence. Swept-carrier time-domain optical data storage also makes use of the spectral interference between a data wave and a reference wave to store the Fourier transform of the temporal data sequence. However, rather than utilizing a fixed carrier frequency for the data, reference, and read waves, in swept carrier, the carrier frequency is varied with time. Note that the theoretical spectral capacity in time-domain and swept-carrier time-domain implementations of spectral data storage is still equal to the ratio of the inhomogeneous linewidth to the homogeneous linewidth ($\Gamma_i/\Gamma_h$).

Full realization of the theoretical spectral capacity using time-domain and swept-carrier implementations has been hindered by two major factors: inability to utilize the entire inhomogeneous linewidth and excitation-induced frequency shifts. Since full utilization of the inhomogeneous linewidth requires high data bandwidths for time-domain implementations or a large laser tuning range for swept-carrier implementations, full bandwidth utilization can be achieved through the use of appropriate support technology. One of the previously stated goals of this project was the development of a frequency agile source to enable achievement of high density storage and extensive effort has been extended to achieve a source that is resonant with the storage material, has a large tuning range and a rapid tuning rate. Overcoming the effects of excitation-induced frequency shifts, however, requires a fundamental change in the data storage process. During the course of ODRAM development we have demonstrated how to overcome the effects of excitation induced frequency shifts, using a new technique, Spatially Distributed Spectral Storage (SDSS).

### 3) Optical Dynamic RAM (ODRAM)

Spectral memory, the ability to store multiple bits in a single location, is the enabling technology behind Optical Dynamic RAM (ODRAM). Spectral storage of multiple bits in a single spatial location allows for non-mechanical access to large stores of data. Utilization of nonmechanical addressing can enable faster data access than is possible with mechanically accessed data storage devices. In particular, latency times can be as short as 10 microseconds, compared to 10 milliseconds for mechanically accessed, disk-based data storage systems. In general, ODRAM will store M bits per spatial location in an array of NxN nonmechanically accessed spatial locations. Projections for initial devices from Phase I were for 10-50 gigabytes of capacity by storing data on a single crystalline slab of storage material.

An ODRAM optical memory system, based on swept-carrier time-domain optical memory is shown in Figure 1 and Figure 2. The overall ODRAM system is depicted in Figure 1. The system consists of an optical unit, an archival backup unit, and an electronic interface and control unit. The archival backup unit consists of a bank of one or more hard drives, which is used to mirror all data in case of power failure. The electronic interface and control unit buffers data to/from the system data bus, performs electronic error correction, and controls a dynamic refresh cycle. This dynamic refresh cycle rewrites data to the storage material during quiescent periods to alleviate decay of the data, which takes place over several hours. Note that silicon DRAM also uses a dynamic refresh cycle, as the data storage time in silicon DRAM is of the order of milliseconds.

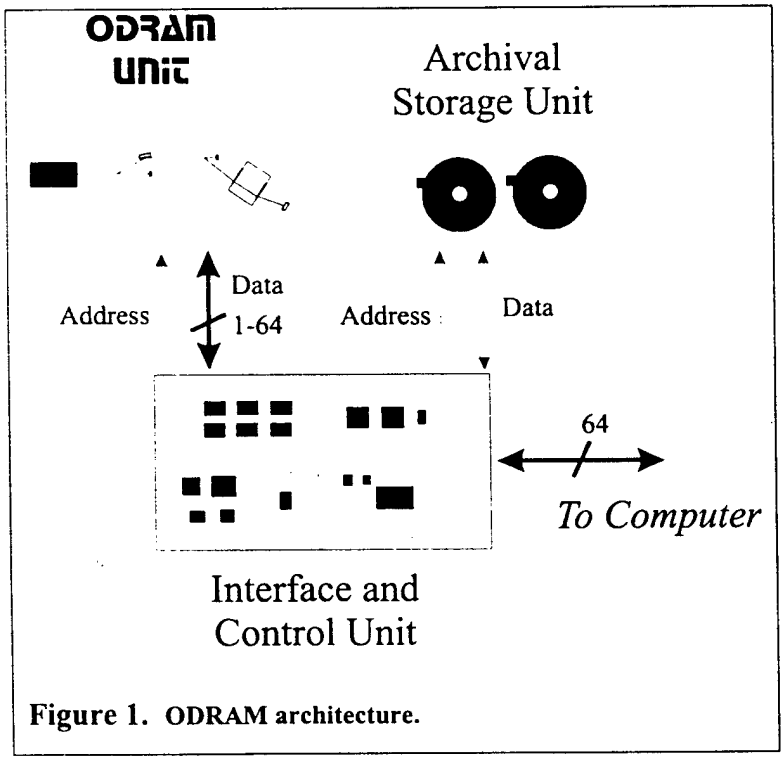


**Figure 1. ODRAM architecture.**

The optical subunit is shown in Figure 2. A frequency-swept optical carrier is generated with a tunable laser. During writing and reading, the carrier frequency is scanned across the inhomogeneous band of the spectral storage material. The constraints on the laser system include a tuning range sufficient to cover the inhomogeneous linewidth of the spectral storage material and a tuning rate sufficient to cover the inhomogenous linewidth within the duration of a single spectral record. The frequency

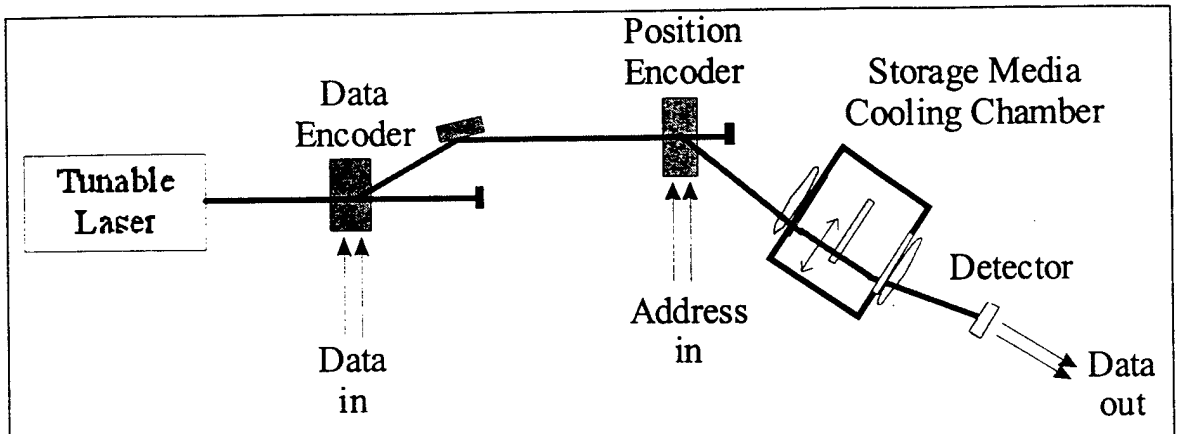


**Figure 2. The optical subunit of ODRAM.**

swept output is passed through the data encoder to generate data and reference beams during the write cycle or a single read beam during the read cycle. The output from the data encoder is passed through an acousto-optic deflector (position encoder) for nonmechanical beam positioning upon the storage material, which is held at the design operating temperature in a cyrostat. During the read cycle, the output signal passes onto a detector and is then electronically processed to yield the desired data.

## Status of Technical Effort:

We have continued the design and development of ODRAM, emphasizing components, such as the storage material, subsystems, such as cryogenics, data modulation units, optical addressing units and the frequency agile laser, and the total system. We have successfully demonstrated a new data storage technique, Spatially Distributed Spectral Storage, that substantially mitigates the effects of excitation induced frequency shifts in spectral storage materials. Utilization of this technique enabled a greater than two order of magnitude increase in spectral capacity than was achievable otherwise. Even more significantly, the results of our demonstrations are extendable to the limit of full utilization of the full inhomogeneous linewidth of the storage material, indicating that single record capacities of greater than 20,000 bits should be readily achievable with our current storage material, $Eu^{3+}:Y_2SiO_5$ and that single record capacities of greater than 50,000 bits might be achievable. Note that at the end of Phase I we projected 100,000 bits per spot or greater, a number used in the Phase II proposal. In addition, we have demonstrated multiple readout of optical data records and started development of dynamic refresh cycles to maintain data integrity.

## Technical Accomplishments

### 1) Overview of demonstration units

Demonstrations of spectral storage were performed using the $^7F_0$-$^5D_0$ transition of site one of $Eu^{3+}:Y_2SiO_5$ ($\lambda$=580.04 nm). An overview of the apparatus is shown in Figure 3. In

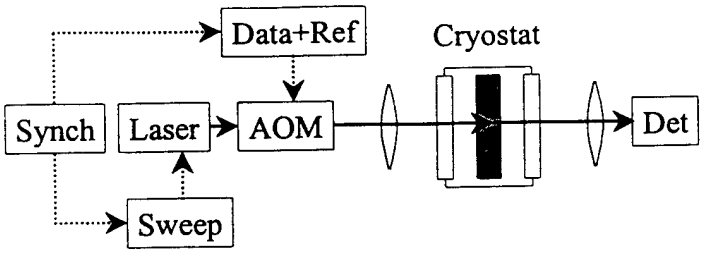


**Figure 3. Overview of apparatus used for single-spot data storage demonstrations. Solid lines indicate optical paths and dotted lines indicate electronic paths.**

brief, a ring dye laser's frequency is locked to an external cavity. The laser frequency is swept sinusoidally with a 6.25 ms period and a sweep range of 1 GHz. The output of the laser is passed through an acoustooptic modulator. During the writing cycle, the acoustooptic modulator is driven with two radio frequency signals thereby generating data and reference beams with a frequency offset of 4 MHz. The data beam is amplitude modulated to encode binary data. During reading, a single radio frequency signal is applied to the acoustooptic modulator to generate the read beam. With the data-reference wave frequency offset utilized here, the data and reference waves emerge from the modulator collinear within the diffraction limit. The output of the acoustooptic modulator passes through a lens and is focussed onto the storage material, a $10\times11\times2$ mm$^3$ crystal of $Eu^{3+}:Y_2SiO_5$ with a doping level of 0.5%, which is held in vacuum at 5K in a flowing He cryostat. Measurements of two-pulse photon echo decays in different spatial regions of the crystal yielded identical time constants within experimental error, indicating negligible temperature gradients across the crystal. The incident radiation propagated along the b axis (2mm dimension) and was polarized along the $D_1$ axis. Following the data storage cycle, illumination of the storage material with the read beam (a duplicate of the reference beam) causes the generation of an outgoing signal that duplicates the data beam. The transmitted read and signal beams interfere, producing a heterodyne beat signal at the data-reference wave offset frequency. The total optical signal is detected and the heterodyne beat signal is electronically extracted and processed to yield the data streams presented below. Note that all data readouts reported here, unless otherwise specified, were performed outside the excited state lifetime of the storage material.

For the Spatially Distributed Spectral Storage (SDSS) demonstrations, the demonstration unit was modified by the inclusion of an acoustooptic deflector, see Figure 4. This acoustooptic deflector was placed in the beam path between the acoustooptic modulator and the lens. The acoustooptic deflector was driven synchronously with the laser frequency sweep, causing the focussed beam to scan across the storage material. For these demonstrations, the peak powers in the data, reference and read beams were all 300 μW. The instantaneous spot size was 40 μm.

Additional electronics were added for the synchronous storage demonstrations, in which the timing of the bits was synchronized to the RF waves driving the acoustooptic modulator. Details of the box labeled "Data+Ref" in Figure 4 for the asynchronous and

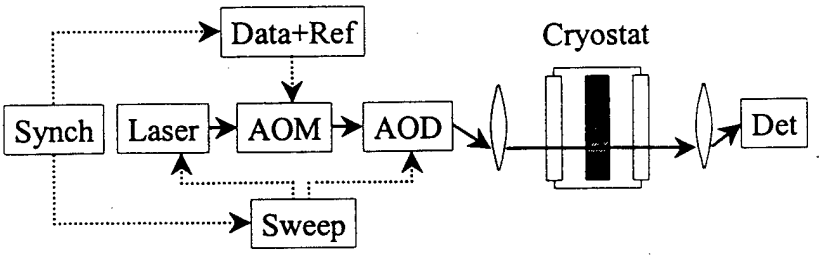


**Figure 4. Overview of apparatus used for spatially distributed spectral storage demonstrations. Solid lines indicate optical paths and dashed lines indicate electronic paths.**

synchronous cases are shown in Figure 5. For the unsynchronized case, electronic data and reference signals were applied to RF switches, imposing the data and reference signals onto RF carriers. These two RF data signals were then combined, amplified, and applied to the acoustooptic modulator. Note that in this case, there is no phase relationship between the RF carriers, the beat frequency, and the data signal. For the synchronous storage demonstrations, the beat frequency between the two RF carriers was used as a clock signal to trigger the data sequence. The advantage of this technique is that a smaller number of oscillations are needed to detect a bit.

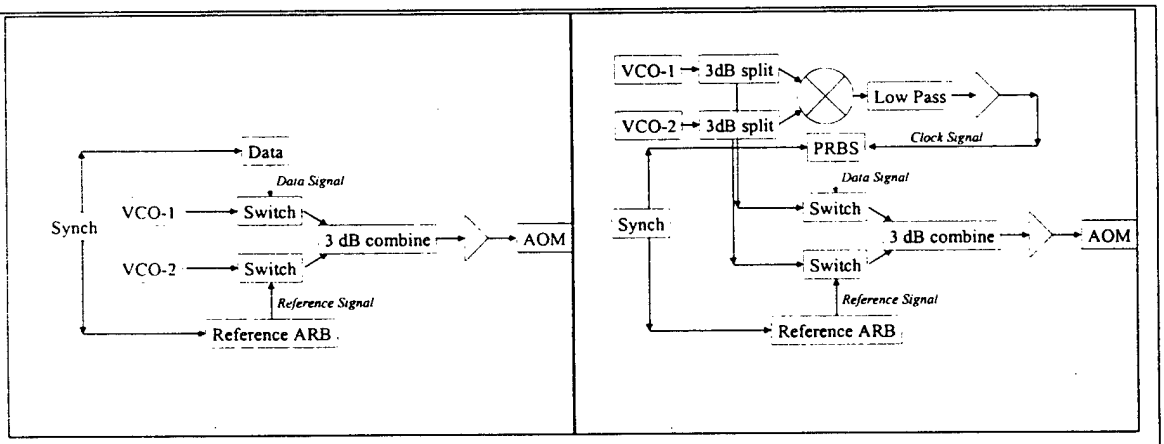


**Figure 5. Depiction of Data+Ref electronics for asynchronous (left) and synchronous data storage demonstrations. In the synchronous demonstrations, a definite phase relationship between the data and the heterodyne beat signal was maintained.**

## 2) Single spot storage results

An example of a small number of bits recalled from a single spatial location is shown in Figure 6. As expected, each "1" is comprised of a 4 MHz beat with a 1 μs envelope. The reconstructed signal efficiency was approximately $10^{-5}$. When attempting to multiplex larger numbers of bits, as in Figure 7, a significant signal degradation was observed for

later bits. A systematic study of the relative durations and starting times of the read, reference, and data windows demonstrated that, for data readout at times longer than the excited state lifetime, the excitation density present prior to the recording of a bit substantially affected the efficiency of readout of that bit. Increasing the excitation density following recording of a bit also affected the readout efficiency for data readout at times longer than the excited state lifetime, but not nearly as strongly. In contrast, for the case of data readout within the excited state lifetime, we observed that the effect of increasing excitation following the recording of a bit had a larger impact on the efficiency than the effect of increasing excitation preceding the recording of a bit. These results indicate that the signal degradation observed in Figure 7 is due to excitation induced frequency shifts.

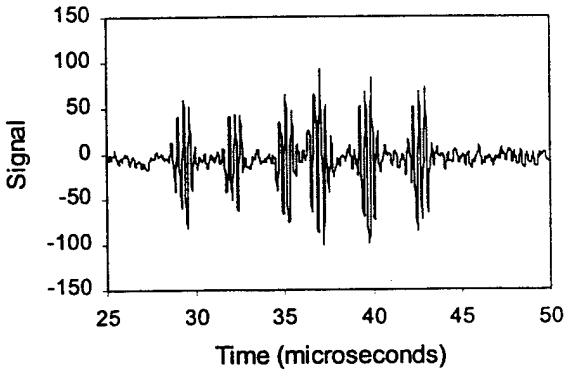


**Figure 6. Observed output signal obtained using heterodyne detected swept carrier data storage under consitions of storage and recall of a small number of bits.**

### 3) Estimating capacity in the presence of excitation induced frequency shifts

As shown in Figure 7, excitation induced frequency shifts can lead to a diminished signal to noise ratio and/or a decreased spectral capacity. In general, excitation-induced frequency shifts lead to a reduced capacity because the effective homogeneous linewidth increases (and $\Gamma_i/\Gamma_h$ thereby decreases) as the number of excited ions in a frequency-selective material increases. As the absorption frequency of any given active ion depends upon its local environment, excitation of one or more ions, which leads to local environmental changes, can cause shifts in the absorption frequencies of other nearby ions. This excitation-induced frequency shift can be manifest as an increase in the effective homogeneous linewidth. Previous experiments on excitation-induced frequency shifts have demonstrated that the broadening depends upon the

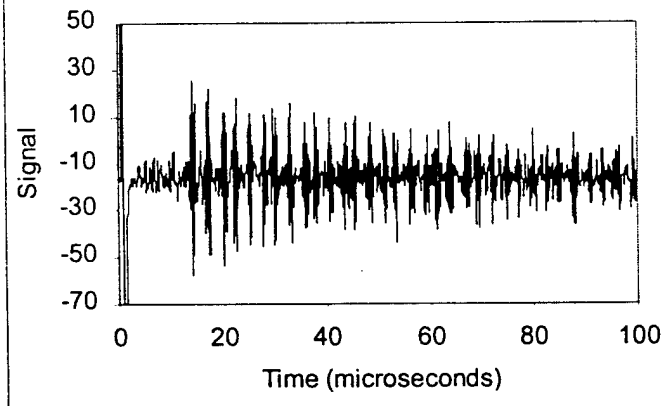


**Figure 7. Observed output signal obtained using heterodyne detected swept carrier data storage under conditions of storage and recall of a large (compared to Figure 6) number of bits. Note that the efficiency of data recall decreases at later times, a signature of excitation induced frequency shifts.**

density of excited ions.[8-10] In practical terms, this means that as one attempts to spectrally multiplex more data by simultaneously exciting more of the inhomogeneous distribution, the spectral capacity of the material decreases, making it impossible to realize the commonly quoted theoretical spectral capacity of a material, which is routinely calculated on the basis of low-excitation level measurements.

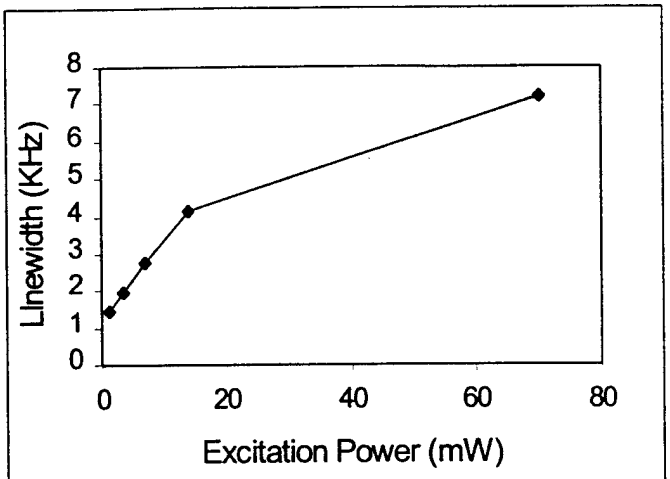


**Figure 8. Measured homogeneous linewidth determined from two pulse echo decays as a function of excitation power for a fixed spot size.**

The experimental demonstrations presented here were performed in a 0.5% europium doped $Y_2SiO_5$ crystal. In the low excitation limit, this crystal has a homogeneous linewidth of approximately 1 kHz and an inhomogeneous width of 10 GHz, indicating a theoretical capacity of $10^7$ bits per spot. However, in agreement with previous studies of excitation induced frequency shifts, we find that homogeneous linewidth obtained from a two-pulse photon echo experiment depends upon the excitation intensity. The results of one such study are shown in Figure 8, which is a plot of measured homogeneous linewidth versus excitation power for a fixed spot size. In this experiment, the excitation bandwidth was approximately 4 MHz. Looking at Figure 8, it is clear that at low excitation intensity, the homogeneous linewidth broadens linearly with increasing excitation intensity. However, at incident power levels above 15 mW, the rate of increase is reduced. We interpret the turnover point as the excitation intensity required to saturate the fraction of the inhomogeneous linewidth of the material spanned by the 4 MHz of excitation bandwidth. For data storage, operation near saturation is desired because the efficiency of data readout is maximized. Upon full excitation of the 4 MHz spanned by the excitation pulses, we find that the linewidth has broadened by about 3 kHz (from 1 kHz to 4 kHz). Therefore, the homogeneous linewidth broadens by approximately 0.8 kHz for each MHz interval of the inhomogeneous distribution that is fully excited. Extrapolating to the simultaneous excitation of the entire 10 GHz wide inhomogeneously broadened line leads to a predicted homogeneous linewidth of approximately 8 MHz and a nearly 4 order of magnitude decrease in the crystal's spectral storage capacity.

### 4) Spatially Distributed Spectral Storage (SDSS)

The reduction in capacity arising from excitation-induced frequency shifts in swept-carrier data storage can be substantially compensated by a new data recording method: Spatially Distributed Spectral Storage (SDSS).[12] SDSS combines spatial scanning with swept-carrier data storage. The laser spot is physically scanned across the storage material as the laser frequency is swept. As a result, the total number of ions excited in

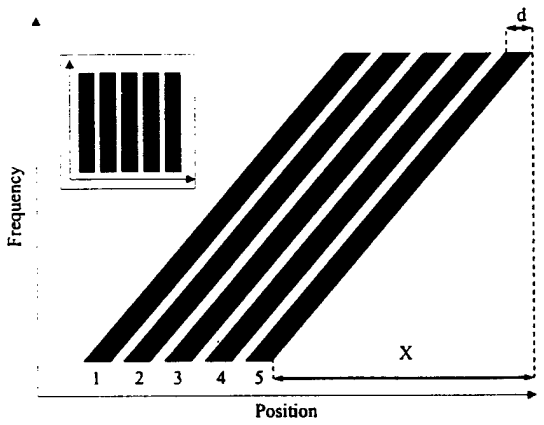


**Figure 9. Schematic representation of the SDSS process. Data records (1-5), represented by bars in the figure, are distributed in position and frequency by simultaneously scanning the excitation frequency and excitation position during writing. Subsequent records can spatially overlap previous records provided that the spectral contents do not overlap. For comparison purposes, the equivalent picture for classic swept-carrier optical memory, in which the position of the excitation beam is constant throughout the frequency sweep and spatially offset for each subsequent record, is given in the inset. The enhancement in capacity, neglecting edge effects, is equal to the ratio of the sweep range to the instantaneous beam diameter, X/d.**

any given spatial storage region is a small fraction of the total number of ions excited during the frequency sweep. Multiple storage sweeps may be applied to each spatial location provided that each sweep utilizes a different portion of the spatial site's spectral bandwidth and that the interval between storage events is longer than the excited state lifetime.

A depiction of how SDSS allows us to fully access the combined spatial and spectral capacity of a storage material is shown in Figure 9. Each bar represents the trajectory through position and frequency of a given data record. For comparison, in classic swept-carrier optical data storage,[5-6,12] wherein the position of the excitation light is fixed during the spectral sweep, the trajectory of a data record appears vertical, rather than slanted, as shown in the inset to Figure 9. The increase in achievable spectral capacity introduced through the use of SDSS is approximately equal to the number of spatial storage locations accessed during the spatial sweep, *i.e.*, the spatial sweep range divided by the instantaneous beam diameter. Therefore, with sufficiently large spatial sweep ranges, orders of magnitude improvement in capacity are achievable. Note that although simple linear data record trajectories through are depicted in Figure 9, more complex trajectories can also be utilized.

Figure 10 shows reconstructed data signals as a function of spatial sweep range. If the laser position is fixed, as in classic swept-carrier data storage (Figure 10a), a small number of bits are reconstructed with a relatively poor signal to noise ratio. As time progresses during the spectral sweep, an increasing number of ions are excited causing a larger effective homogeneous linewidth and a reduced reconstruction efficiency. For short spatial sweeps, Figure 10b and Figure 10c, the excited ion density is reduced (relative to Figure 10a) and improvements in reconstruction are observed. However, the reconstructed signal intensity still decays at later times. In Figure 10d, the spatial sweep extends across 120 beam diameters (over 3.25 ms) and the reconstructed data has a relatively time-independent amplitude, indicating that excitation-induced frequency shifts have been eliminated as a major relaxation factor. Figure 11 shows the complete reconstruction (of which Figure 3d is a subset) of a 2000 bit sequence over 2 ms. As predicted, the effective storage capacity in a record has been increased from the spatially

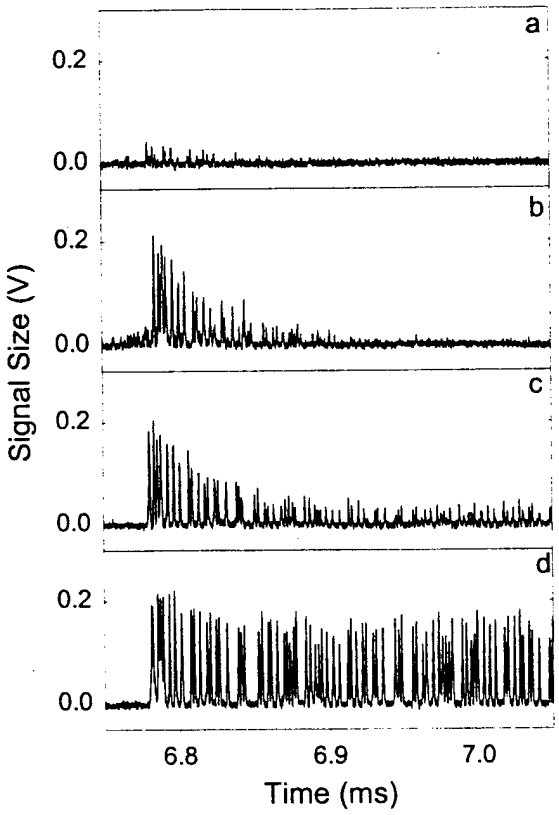


**Figure 10.** **Data reconstruction for four different spatial sweep ranges. The bit duration was 1 μs, the offset frequency between data and reference beams was 4 MHz, and the instantaneous spot size was 40 μm. The spatial ranges covered during the 6.75 to 7.05 ms time window depicted in this figure were a) 0μm, b) 36 μm, c) 160 μm, d) 440 μm, corresponding to average sweep rates of a) 0 μm/ms, b) 120 μm/ms, c) 530 μm/ms, d) 1500 μm/ms.**

fixed situation by a factor close to the number of individual spatial locations addressed (120). In a separate experiment with a slightly higher bit rate (not shown), we achieved storage and recall of 3300 bits. It should be noted that due to the sinusoidal nature of the spatial and spectral sweeps, the scan rates are not constant.

There are limits to the improvement in storage capacity achievable with SDSS. Some obvious limitations include the homogeneous linewidth in the low excitation limit and incomplete spectral utilization at the physical edges of the storage material. Note that the impact of this second limitation can be minimized by more advanced design of the data record trajectories compared to the simple linear trajectories shown in Figure 9. More importantly, signal fidelity can be degraded for very high spatial sweep rates, due to poor spatial and spectral overlap between the data and reference beams. In particular, the spatial displacement rate has to be slow enough that the spectral content deposited by the reference beam in a given spatial location substantially matches the spectral content deposited by the data beam in the same spatial location. In other words, the time required to sweep across a spatial location should be larger than the time required to sweep the spectral content of both beams through the same spectral region. For the current situation, wherein the reference and read beams are collinear and copropagating, this requirement can be expressed as:

$$\frac{d}{r_{spatial}} > \frac{\Delta\upsilon + 1/(2\tau_B)}{r_{spectral}} \qquad (1)$$

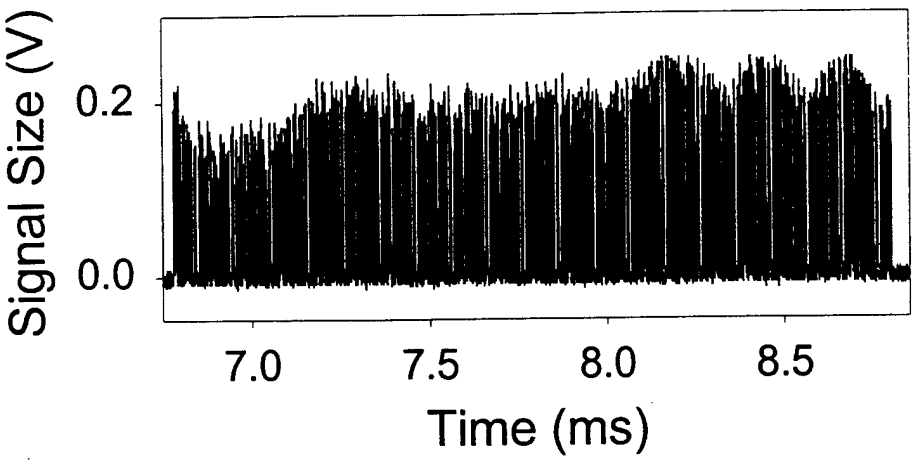


**Figure 11. Data reconstruction of a 2000 bit sequence. The bit duration was 1 μs and the offset frequency between data and reference beams was 4 MHz. The total spatial sweep range was 4800 μm, of which, approximately 4100 μm was used to store data.**

Here, $d$ is the diameter of the beam, $\Delta\upsilon$ is the frequency offset between the data and reference beams, $\tau_B$ is the bit duration, $r_{spatial}$ is the spatial sweep rate and $r_{spectral}$ is the spectral sweep rate. At the maximal sweep rates achieved, in the central region of Figure 11, $d/r_{spatial}$ was 8.3 μs and $(\Delta\upsilon+1/(2\tau_B))/r_{spectral}$ was 4.5 μs. These numbers indicate that in future development we would be able to employ higher spatial sweep rates to offset the effects of higher input powers, which are needed (at constant spot size) for higher data rates. Additional improvement could be obtained by spatially offsetting the data and reference beams to account for their frequency offset.

### 5) Demonstration of multiple readouts of a single record

The readout process used in ODRAM is inherently destructive, resulting in a signal degradation with each successive read of a given data record. We have characterized the degradation in signal fidelity resulting from reading out a given data record multiple times. Example data is shown in Figure 12. As can be seen, for there is a large decrease in signal size and fidelity during the first couple of readouts. Signal size and fidelity continues to decrease with subsequent readouts, although at a slower rate. According to these results, it should be possible to read a record between three and five times before refresh would be required.

### 6) Development of techniques for data refresh

One of the difficulties associated with dynamic refresh cycles in spectral data storage is a long persistence of data in the storage material that has degraded to the point that it is not reliably readable, but still has a large enough amplitude to significantly interfere with, and thereby affect the fidelity of, any new data that is stored in the same spatial/spectral location. In general, this implies that prior to writing data in a previously exposed spot,

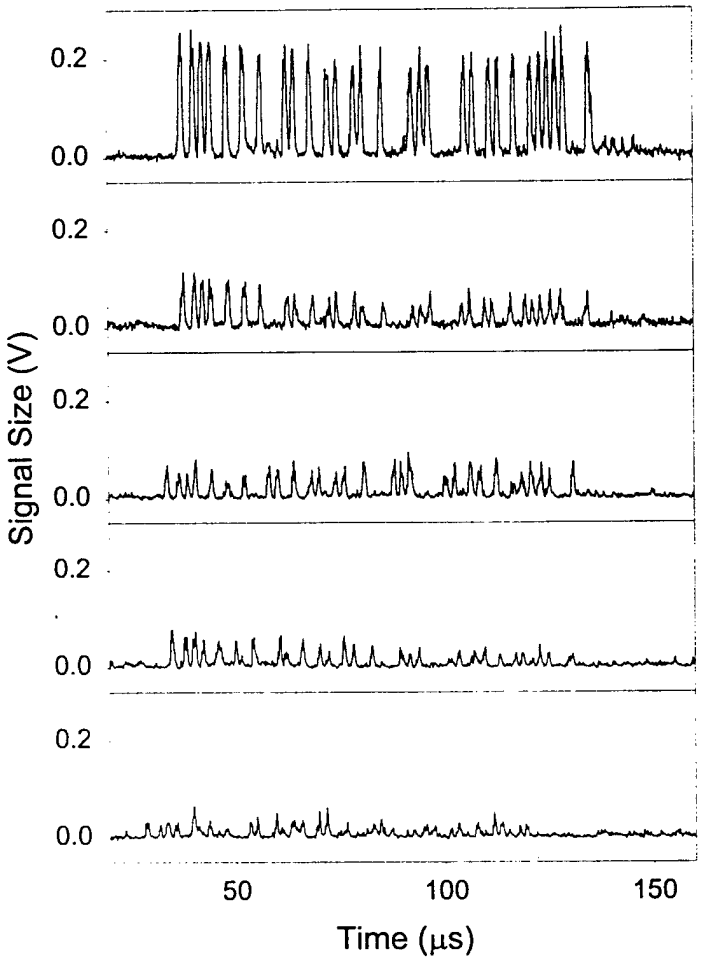


**Figure 12. Data readout signal as a function of the number of read cycles. Starting from the top, the number of read cycles is one, three, five, seven, and nine.**

complete erasure of the old data is necessary. In the case of dynamic refresh, this is also the case unless the new data can be inscribed upon the storage material such that the new spatial/temporal grating constructively interferes with the preexisting spatial/temporal grating. We have begun development of a technique for locking the phases of subsequent writing cycles by locking the data sequence to the phase of the heterodyne beat signal. The preliminary electronics that accomplish this task are shown in Figure 5. This technique also has a second advantage, a higher bit rate given a specific delay time between when the reference beam interrogates a specific frequency channel and when the data beam interrogates the same frequency channel.

## Comparison to Competing Technologies and the Future of ODRAM

Data storage is a key component of all modern computing systems. Achievable computational performance strongly reflects limitations inherent in the storage technologies currently available. Crucial performance characteristics of commercially viable data storage devices include:

- random access latency time
- peak data transfer rate
- total capacity
- persistence time
- cost per unit performance.

In current desktop, server and high performance computing applications, overall data storage functionality is achieved through a combination of technologies, wherein the strengths of one technology compensate for the weaknesses of others. This situation has arisen because no single storage technology can provide optimal performance in all arenas. Two memory technologies currently dominate high performance storage applications: semiconductor RAM and mechanically accessed magnetic devices. Other technologies are important for archival backup (tape) and distribution (CD-ROM). Semiconductor RAM is very fast in terms of transfer rate (multi-Gigabit/second bus aggregate) and random access time (10 nanoseconds). However, it is expensive and usually volatile. Mechanically accessed magnetic storage devices offer enormous (up to 30 Gigabyte/unit) permanent storage for a very low price, but have very long (8-10 millisecond) random access latency times and relatively low peak transfer rates (100-300 megabit/second serial).

The mismatch in performance (speed and capacity) between magnetic storage devices and semiconductor RAM leads to performance bottlenecks in many computational applications, especially those applications that require rapid access to large numbers of small data records. It is exactly this bottleneck that ODRAM was designed to address. In a simple implementation, ODRAM is a two dimensional memory with data stored in an array of spatial locations. At first glance, this seems conceptually similar to standard magnetic or optical disk technology. However, in ODRAM, many thousands of bits are stored in each and every spatial location, through the techniques of persistent spectral hole burning and spectral holography. Because of the high multiplicity of bits within a single spatial address, substantial data capacities can be achieved with a relatively small number of spatial addresses. Therefore, multi Gigabyte (10-50) capacity storage units can be implemented with entirely non-mechanical beam steering- thus providing random access times of microseconds rather than the random access times of milliseconds common in mechanically accessed magnetic disk storage technology.

Although ODRAM offers distinct performance advantages relative to mechanically accessed mass storage devices, its performance, as measured by latency time and peak transfer rate does not match semiconductor DRAM. Therefore, for ODRAM to become commercially viable, the price point must lie below the price point of semiconductor DRAM. The upper solid curve in Figure 13 is a projection of

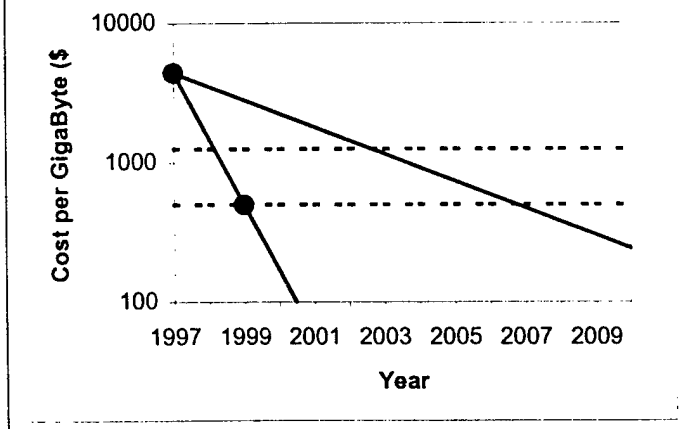


**Figure 13. Projected DRAM pricing (upper solid line), actual DRAM pricing (lower solid line with dots), projected price for a 20 gigabyte ODRAM (upper dashed line) and projected pricing for a 50 gigabyte ODRAM.**

semiconductor DRAM pricing based on historical trends up to 1997. According to that historical projection, which we made during our Phase I analysis of ODRAM, the break even price point between a 50 Gigabyte ODRAM, the lower dashed curve in Figure 13 and 50 gigabytes of semiconductor DRAM should occur in 2007. However, over the past two years, the price of silicon DRAM has fallen much at a much faster rate than would be predicted from historical trends up to 1997, as depicted by the lower solid curve in Figure 13. Note that the crossing point in pricing between silicon DRAM and ODRAM already now been reached, almost 10 years earlier than historical projections would have indicated. Given the large number of expensive subsystems (cryogenics, the frequency agile laser, acoustooptic deflector, *etc.*) necessary for the implementation of ODRAM, achieving competitive price points by innovation and redesign will be extremely difficult if not impossible.

## References


1. R. Equall, Y. Sun, R. Cone and R. Macfarlane, Phys. Rev. Lett., **72**, 2179 (1994).
2. R. Yano, M. Mitsunaga and N. Uesugi, J. Opt. Soc. Am. B, **9**, 992 (1992).
3. W.R. Babbitt and T.W. Mossberg, Opt. Comm., **65**, 185 (1988).
4. Y.S. Bai and R. Kachru, Opt. Lett., **18**, 1189 (1994).
5. H. Lin, T. Wang, G.A. Wilson and T.W. Mossberg, Opt. Lett., **20**, 282 (1995).
6. H. Lin, T. Wang and T.W. Mossberg, Opt. Lett., **20**, 1528 (1995).
7. X.A. Shen and R. Kachru, Appl. Opt., **36**, 6692 (1997).
8. H. Lin, T. Wang, G.A. Wilson and T.W. Mossberg, Opt. Lett., **20**, 928 (1995)
9. J. Huang, J.M. Zhang, A. Lezama and T.W. Mossberg, Phys. Rev. Lett., **63**, 78 (1989).
10. J. Huang, J. Zhang and T.W. Mossberg, Opt. Comm., **75**, 29 (1990).
11. M. Mitsunaga, T. Takagahara, R. Yano and N. Uesugi, Phys. Rev. Lett., **68**, 3216 (1992).
12. A.E. Johnson, E.S. Maniloff and T.W. Mossberg, Opt. Lett., **24**, 1526 (1999).
13. T.W. Mossberg, Opt. Lett., **17**, 535 (1992).

## Personnel Supported

Alan E. Johnson

Eric S. Maniloff

Thomas W. Mossberg

## Publications:

1) A.E. Johnson, E.S. Maniloff and T.W. Mossberg, "Spatially Distributed Spectral Storage", Optics Letters, **24**, 1526 (1999).

2) A.E. Johnson, E.S. Maniloff and T.W. Mossberg, "Implementation of Optical Dynamic RAM using Spatially Distributed Spectral Storage", Proceedings of SPIE, **3802**, 112 (1999).

3) E.S. Maniloff, A.E. Johnson and T.W. Mossberg, "Spectral Data Storage Using Rare-Earth Doped Crystals", MRS Bulletin, **24**, 46 (1999).

## Interactions/Transitions:

1) A.E. Johnson, E.S. Maniloff and T.W. Mossberg, "Optical Dynamic RAM: A commercially viable implementation of Spectral Data Storage?" presented at SPIE Annual Meeting, July 1999.

2) E.S. Maniloff, A.E. Johnson and T.W. Mossberg, "Spatially Distributed Spectral Storage", presented at Workshop on Applications of Spectral Hole Burning, March 1999.

3) E.S. Maniloff, A.E. Johnson and T.W. Mossberg, "Spectral Data Storage: Technical Possibilities and Challenges" presented at National Storage Industry Consortium annual meeting, June 1999.

## New discoveries, inventions, or patent disclosures:

1. Method of mitigating the effects of excitation induced frequency shifts in spectrally selective materials. Provisional patent application filed with the US patent office 10/98. U.S. utility application number 09/412,841and PCT application number US99/23236 filed 10/5/99.